arXiv:1702.05548v1 [math.OC] 18 Feb 2017

Bi-Level Online Control without Regret

Andrey Bernstein *

February 21, 2017

Abstract

This paper considers a bi-level discrete-time control framework with real-time constraints, consisting of several local controllers and a central controller. The objective is to bridge the gap between the online convex optimization and real-time control literature by proposing an online control algorithm with small dynamic regret, which is a natural performance criterion in nonstationary environments related to real-time control problems. We illustrate how the proposed algorithm can be applied to real-time control of power setpoints in an electrical grid.

1 Introduction

We consider a bi-level discrete-time control framework with real-time constraints, consisting of several *local controllers* and a *central controller*. The tasks of a local controller are (i) to implement setpoints issued by the central controller and (ii) to advertise a prediction of the objective function and the constraints on the feasible setpoints to the central controller. In turn, the central controller uses these advertisements and its system-wide measurements and modeling to compute the next feasible setpoints for the local controllers. This framework is illustrated in Figure 1.

This framework is appropriate in the modern real-time control of cyber-physical systems, such as electrical-grid control (e.g., [20, 22, 11, 5, 12]) or the control of autonomous vehicles (e.g., [23, 24, 8]). In particular, in the context of electrical-grid control, a similar approach was used in [5, 25] to control the power setpoints of devices in real time. Moreover, [12] considered a distributed control framework in this spirit. However, in [5, 25], a heurisitic was used, and no theoretical guarantees were provided, whereas [12] assumed perfect knowledge of the objective function and feasible sets at the time of the decision making at the central controller. In [4], an approach for error correction in local controllers was proposed; however, the performance of the closed-loop system was not analyzed theoretically.

Observe that this framework is reminiscent of online learning and, in particular, of online convex optimization (OCO) algorithms. Indeed, at each time step, the central controller is faced with a nonstationary (time-varying) optimization problem, and it is required to track its optimal solution. The literature on OCO is focused on devising online algorithms that have provably bounded (or vanishing) *average regret* [6, 26, 9, 17]; see also [16] for a recent overview of the subject. However, these algorithms typically lack the control perspective because they naturally operate in an *open-loop* fashion. In particular, an online optimization algorithm issues setpoints that are assumed to

*This work was supported by the U.S. Department of Energy under Contract No. DE-AC36-08GO28308 with the National Renewable Energy Laboratory. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

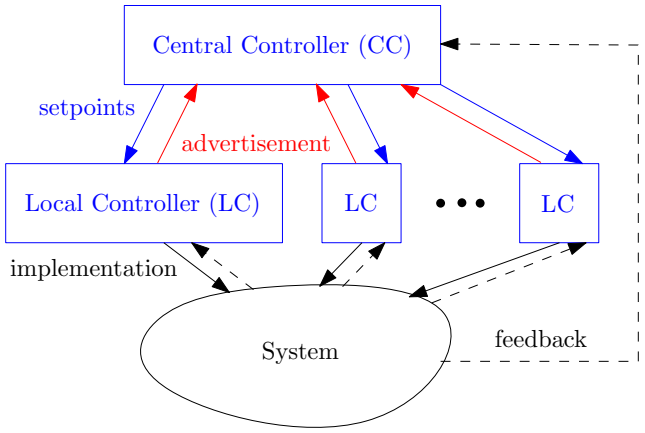

Figure 1: Illustration of the bi-level control framework.

be "implemented" perfectly by the system. Hence, contrary to typical control settings, there is no explicit feedback from the system on the actual implementation.

This paper is the first attempt to bridge the gap between online convex optimization and real-time control. We propose a first-order online control algorithm in the spirit of [26], which we call the Online Gradient Control (OGC), in the above-mentioned bi-level framework. We show theoretical guarantees on the algorithm's *dynamic regret* (or *tracking regret*). The latter is an extension of the standard regret notion, which compares the performance of an online algorithm to an arbitrary sequences of setpoints (rather than a single fixed setpoint) in hindsight (see, e.g., [19, 1, 18, 15]). This regret notion is natural in nonstationary environments associated with real-time control problems. We also discuss the application of the framework and algorithm to the real-time control of heterogeneous devices and, in particular, to the real-time control of power setpoints in an electrical grid.

An additional contribution of this paper is in the context of online convex optimization with *time-varying feasible sets*. Indeed, as the online control problem involves sets of feasible setpoints that change with time, the no-regret result established in this paper is also applicable in the "open-loop" optimization setting. To the best of our knowledge, the only work that considers time-varying feasible sets explicitly is that of [21]; however, it is assumed there that these sets are drawn from a *fixed* unknown distribution, whereas in the present paper we assume an *arbitrary* sequence of feasible sets.

The paper is organized as follows. Section 2 presents the bi-level control framework and relevant notation. Section 3 introduces our OGC algorithm and analyzes its dynamic regret. Section 4 shows how to apply the algorithm to control a mix of heterogeneous resources, and, in particular, to control an electrical grid. Finally, Section 5 concludes the paper and outlines some future research directions.

2 Bi-Level Control Framework

Assume that there are J local controllers (LCs) indexed by $j = 1, \ldots, J$. The discrete time step index is denoted by $n = 1, 2, \ldots$. Let $\mathcal{S}_n(j)$ denote the convex compact set of feasible setpoints of LC j at time step n. Also, let $C_n^{(j)} : \mathcal{S}_n(j) \to \mathbb{R}$ denote a convex cost function that represents the

objective function of LC j. At each time step n, controller j sends to the central controller (CC) an *advertisement* $(\mathcal{A}_{n+1}(j), \hat{C}_{n+1}^{(j)})$ of its feasible set and cost function valid for time step $n+1$ by using *a persistent predictor*, namely $\mathcal{A}_{n+1}(j) = \mathcal{S}_n(j)$ and $\hat{C}_{n+1}^{(j)} = C_n^{(j)}$.

Upon receiving the advertisements from all the LCs, the CC computes the feasibility constraints on the overall system based on its system view and the advertisements. Let $\mathcal{U}_n \subseteq \mathcal{A}_{n+1} := \mathcal{S}_n(1) \times ... \times \mathcal{S}_n(J)$ denote the compact convex set representing the system feasibility constraints. The CC also computes the estimation of the overall objective function:

$$\begin{aligned} F_n(x) &:= \sum_{j=1}^{J} w_j \hat{C}_{n+1}(x(j)) + G_n(x) \\ &= \sum_{j=1}^{J} w_j C_n(x(j)) + G_n(x) \end{aligned} \tag{1}$$

for any $x \in \mathcal{U}_n$, where w_j are some weighting (normalization) factors, and the convex function $G_n(x)$ represents a system-wide objective. Finally, the CC computes a vector of setpoints x_{n+1} based on $F_n(x)$ and $\mathcal{U}_n$, and sends the individual setpoints $x_{n+1}(j)$ to the LCs. Upon receiving $x_{n+1}(j)$, LC j implements a feasible approximation $y_{n+1}(j) \in \mathcal{S}_{n+1}(j)$, and the process repeats. The interaction between the LCs and the CC is summarized in Algorithm 1.

Algorithm 1 Interaction between local controllers and central controller

1: Set $n = 0$.
2: **loop**
3: At time step n, every LC $j = 1, ..., J$:
   (a) Receives a setpoint request $x_n(j)$ sent by the CC.
   (b) Implements an approximation $y_n(j)$ of $x_n(j)$. The implemented setpoint $y_n(j)$ is constrained to lie in the set $\mathcal{S}_n(j)$ representing the local feasibility constraints.
   (c) Predicts its feasible set $\mathcal{S}_{n+1}(j)$ by $\mathcal{A}_{n+1}(j) := \mathcal{S}_n(j)$.
   (d) Predicts its local objective function $C_{n+1}^{(j)}$ by $\hat{C}_{n+1}^{(j)} := C_n^{(j)}$.
   (e) Sends to the CC $\mathcal{A}_{n+1}(j)$ and $\hat{C}_{n+1}^{(j)}$ over a communication network.
4: Upon receiving the advertisements from all the LCs, the CC:
   (a) Computes the feasibility constraints on the overall system $\mathcal{U}_n \subseteq \mathcal{A}_{n+1} := \mathcal{S}_n(1) \times ... \times \mathcal{S}_n(J)$.
   (b) Computes the setpoints' vector $x_{n+1} \in \mathcal{U}_n$ based on its current objective function (1).
   (c) Sends $x_{n+1}(j)$ to every LC j over a communication network.
5: $n := n + 1$.
6: **end loop**

3 Online Gradient Control

In this section, we present our Online Gradient Control (OGC) algorithm. The algorithm is based on the following two steps:

(i) The central controller chooses setpoints according to the online gradient descent algorithm:

$$x_{n+1} = \mathcal{P}_{\mathcal{U}_n}(\hat{y}_n - \alpha \nabla F_n(\hat{y}_n)), \tag{2}$$

where $\mathcal{P}_{\mathcal{U}}(\cdot)$ is the projection operator, α is a step-size parameter, and $\hat{y}_n$ is the estimation of the setpoint y_n implemented by the LCs at time step n. The latter is obtained either from the LCs or by system-wide measurements.

(ii) Upon receiving a setpoint $x_n(j)$, LC j implements a projected version thereof, namely:

$$y_n(j) = \mathcal{P}_{\mathcal{S}_n(j)}(x_n(j)). \tag{3}$$

We next analyze the performance of the OGC in terms of optimality and stability. To this end, we first introduce the following assumptions and definitions.

Assumption 1. *The measurement of the implemented setpoint $\hat{y}_n$ is ε-accurate. That is, for all n,*

$$\|y_n - \hat{y}_n\| \leq \varepsilon.$$

Assumption 2. *The gradients $\{\nabla F_n\}$ are uniformly bounded and Lipschitz continuous with a common parameter $\lambda < \infty$. Namely, for all n, all $x, x' \in \mathcal{U}_n$, $\|\nabla F_n(x) - \nabla F_n(x')\| \leq \lambda \|x - x'\|$. Let F denote a finite constant that is a uniform upper bound on $\|\nabla F_n(x)\|$.*

Assumption 3. *The sequence of feasible sets $\{\mathcal{S}_n\}$ is uniformly bounded. That is, there exists a finite constant B such that $|\mathcal{S}_n| \leq B$ for all n. Let D denote the upper bound on the diameters of $\{\mathcal{S}_n\}$, so that $\mathsf{diam}(\mathcal{S}_n) \leq D$ for all n.*

Definition 1 (Admissible sequence of setpoints)**.** A sequence $\mathbf{x}_{1:n} := \{x_k\}_{k=1}^n$ is said to be *admissible* if $x_k \in \mathcal{U}_k$ for every $k = 1, ..., n$. Let $\mathcal{X}_n$ denote the set of all the admissible sequences of length n, and let $\mathcal{X}_\infty$ denote the set of all the infinite admissible sequences $\mathbf{x}_{1:\infty}$.

Definition 2 (Dynamic regret)**.** Consider the sequence of implemented setpoint $\mathbf{y}_{1:n} := \{y_k\}_{k=1}^n$ by the control algorithm up to time step n. The *total dynamic regret* of the algorithm with respect to a given sequence $\mathbf{z}_{1:n} \in \mathcal{X}_n$ is defined as

$$r_n(\mathbf{z}_{1:n}) := \sum_{k=1}^{n} (F_k(y_k) - F_k(z_k)). \tag{4}$$

Similarly, $r_n(\mathbf{z}_{1:n})/n$ is the *average dynamic regret*.

Definition 3 (Temporal variability)**.** For any sequence $\mathbf{x}_{1:n} = \{x_k\}_{k=1}^n$, let

$$V(\mathbf{x}_{1:n}) := \sum_{k=1}^{n} \|x_k - x_{k+1}\| \tag{5}$$

denote its temporal variability.

Theorem 1. *Under Assumptions 1, 2, and 3, for any $\mathbf{z}_{1:\infty} \in \mathcal{X}_\infty$ and any $\alpha > 0$, we have that*

$$\limsup_{n\to\infty} \frac{r_n(\mathbf{z}_{1:n})}{n} \leq K_1\alpha + \frac{K_2(1+\alpha\lambda)}{\alpha}\varepsilon + \frac{K_3}{\alpha}\limsup_{n\to\infty}\frac{V(\mathbf{z}_{1:n})}{n} \tag{6}$$

where

$$K_1 := \frac{F^2}{2},\ K_2 := \frac{[2(D+\alpha F)+(1+\alpha\lambda)\varepsilon]}{2},\ \textit{and}\ K_3 := D + B.$$

In particular, (6) *is valid for the optimal sequence*

$$\mathbf{z}^*_{1:n} \in \operatorname*{arg\,min}_{\mathbf{z}_{1:n}\in\mathcal{X}_n} \sum_{k=1}^{n} F_k(z_k).$$

Proof. The proof follows that of [26, Theorem 1]. For simplicity of exposition, we use below a scalar-style notation; however, the proof works for vectors by interpreting the regular multiplication as the inner product.

Let $\mathbf{z}_{1:n+1} \in \mathcal{X}_{n+1}$. We have that

$$\begin{aligned}
&(y_{n+1} - z_{n+1})^2 \leq (x_{n+1} - z_{n+1})^2 \\
&= (x_{n+1} - z_n + z_n - z_{n+1})^2 \\
&= (x_{n+1} - z_n)^2 + 2(x_{n+1} - z_n)(z_n - z_{n+1}) + (z_n - z_{n+1})^2 \\
&\leq (\hat{y}_n - z_n - \alpha\nabla F_n(\hat{y}_n))^2 \\
&\quad + [2(x_{n+1} - z_n) + (z_n - z_{n+1})](z_n - z_{n+1}) \\
&\leq (\hat{y}_n - z_n - \alpha\nabla F_n(\hat{y}_n))^2 + 2(D+B)\|z_n - z_{n+1}\|,
\end{aligned} \tag{7}$$

where the first inequality follows by using (3), the fact that $z_{n+1} \in \mathcal{U}_{n+1} \subseteq \mathcal{S}_{n+1}$, and the non-expansive property of the projection operator; the second inequality holds by (2), the fact that $z_n \in \mathcal{U}_n$, and the non-expansive property of the projection operator; and in the last inequality, we used the Cauchy-Schwarz inequality and the fact that under Assumption 3

$$\begin{aligned}
\|2(x_{n+1} - z_n) + (z_n - z_{n+1})\| &\leq 2\|x_{n+1} - z_n\| + \|z_n\| + \|z_{n+1}\| \\
&\leq 2\mathsf{diam}(\mathcal{S}_n) + |\mathcal{S}_n| + |\mathcal{S}_{n+1}| \\
&\leq 2(D+B).
\end{aligned}$$

We now expand the first term in (7). It holds that

$$\begin{aligned}
&(\hat{y}_n - z_n - \alpha\nabla F_n(\hat{y}_n))^2 \\
&= (y_n - z_n - \alpha\nabla F_n(y_n) + (\hat{y}_n - y_n) + \alpha(\nabla F_n(\hat{y}_n) - \nabla F_n(y_n)))^2.
\end{aligned} \tag{8}$$

Let

$$\gamma_n := (\hat{y}_n - y_n) + \alpha(\nabla F_n(\hat{y}_n) - \nabla F_n(y_n))$$

and note that under Assumptions 1 and 2, we have

$$\|\gamma_n\| \leq (1+\alpha\lambda)\varepsilon.$$

Continuing the derivation in (8), we obtain

$$\begin{aligned}
&(\hat{y}_n - z_n - \alpha\nabla F_n(\hat{y}_n))^2 \\
&\leq (y_n - z_n - \alpha\nabla F_n(y_n))^2 \\
&\quad + [2(y_n - z_n - \alpha\nabla F_n(y_n)) + \gamma_n]\,\gamma_n \\
&\leq (y_n - z_n - \alpha\nabla F_n(y_n))^2 \\
&\quad + [2(D + \alpha F) + (1 + \alpha\lambda)\varepsilon](1 + \alpha\lambda)\varepsilon \\
&= (y_n - z_n)^2 - 2\alpha\nabla F_n(y_n)(y_n - z_n) + \alpha^2(\nabla F_n(y_n))^2 \\
&\quad + [2(D + \alpha F) + (1 + \alpha\lambda)\varepsilon](1 + \alpha\lambda)\varepsilon \\
&\leq (y_n - z_n)^2 - 2\alpha(F_n(y_n) - F_n(z_n)) + F^2\alpha^2 \\
&\quad + [2(D + \alpha F) + (1 + \alpha\lambda)\varepsilon](1 + \alpha\lambda)\varepsilon
\end{aligned} \tag{9}$$

where the second inequality holds by the Cauchy-Schwarz inequality and Assumptions 2 and 3; and the last inequality holds by the standard argument for comparing the instantaneous regret of linear and strictly convex functions (see, e.g., [26]). Combining (7) and (9) yields

$$\begin{aligned}
&(y_{n+1} - z_{n+1})^2 \\
&\leq (y_n - z_n)^2 - 2\alpha(F_n(y_n) - F_n(z_n)) + F^2\alpha^2 \\
&\quad + [2(D + \alpha F) + (1 + \alpha\lambda)\varepsilon](1 + \alpha\lambda)\varepsilon \\
&\quad + 2(D + B)\|z_n - z_{n+1}\|
\end{aligned} \tag{10}$$

By rearranging (10), we obtain

$$\begin{aligned}
F_n(y_n) - F_n(z_n) &\leq [(y_n - z_n)^2 - (y_{n+1} - z_{n+1})^2]/(2\alpha) + \alpha F^2/2 \\
&\quad + [2(D + \alpha F) + (1 + \alpha\lambda)\varepsilon](1 + \alpha\lambda)\varepsilon/(2\alpha) + (D + B)\|z_n - z_{n+1}\|/\alpha.
\end{aligned}$$

Averaging the last inequality yields

$$\begin{aligned}
\frac{1}{n}\sum_{k=1}^{n}(F_k(y_k) - F_k(z_k)) \quad \leq \quad & [(y_1 - z_1)^2 - (y_{n+1} - z_{n+1})^2]/(2\alpha n) \\
& + K_1\alpha + \frac{K_2(1 + \alpha\lambda)}{\alpha}\varepsilon + (K_3/\alpha)\frac{1}{n}V(\mathbf{z}_{1:n}),
\end{aligned}$$

which completes the proof. □

Observe that Theorem 1 establishes that if the measurement error ε is small, and if the optimal trajectory $\mathbf{z}^*_{1:n}$ varies slowly (or infrequently), in the sense that

$$\limsup_{n\to\infty} \frac{V(\mathbf{z}^*_{1:n})}{n}$$

is small, then the corresponding average dynamic regret will be small for the appropriately chosen step-size α.

Remark 1. The discrete-time control problem considered here is an approximation to the corresponding continuous-time optimal control problem. The result of Theorem 1 thus implicitly states that if the optimal continuous-time trajectory $z^*(t)$ is continuous in t, then one can choose a fine enough discretization of the timescale so that the average time variability $V(\mathbf{z}^*_{1:n})/n$ is asymptotically small, hence yielding small asymptotic regret.

Remark 2. Note that when $\hat{y}_n = x_n$ and the feasible set $\mathcal{U}_n$ does not depend on n, the algorithm (2) is the well-known online gradient descent algorithm first introduced in [26]. The case when $\hat{y}_n = x_n$ but $\mathcal{U}_n$ *depends* on n can be considered as an (open-loop) online optimization setting with time-varying feasible sets, and the result of Theorem 1 applies to this case as well.

We conclude this section by noting that the OGC algorithm is stable in the input-to-state stability sense by construction. To that end, consider the nonlinear dynamical system for the state variable y_n given by

$$y_{n+1} = \mathcal{P}_{\mathcal{S}_{n+1}}(\mathcal{P}_{\mathcal{U}_n}(y_n - \alpha \nabla F_n(y_n) + \epsilon_n)), \tag{11}$$

where ϵ_n is the measurement error associated with Assumptions 1 and 2. Here, the sets $\{\mathcal{S}_n\}$ (and the sequence $\{\epsilon_n\}$) can be considered as *exogenous inputs* to this dynamical system. Under Assumption 3, it is clear that $\|y_n\| \leq B$ for all n, which establishes a bounded-input-bounded-state (BIBS) stability. Indeed, it states that if the sequence of "inputs" $\{\mathcal{S}_n\}$ is uniformly bounded, so is the sequence of "states" $\{y_n\}$.

4 Application to Real-Time Control of Heterogeneous Devices

In this section, we consider the setting where the LCs control heterogeneous devices of two general types: (i) *convex devices* with convex feasible sets and (ii) *discrete* devices with discrete feasible sets with a *finite* number of elements. Observe that for convex devices, the OGC can be directly applied. For discrete devices, we propose the following randomized scheme in the spirit of repeated games. In Section 4.1, we give the general algorithm, whereas in Section 4.2, we outline the application in the context of electrical-grid control.

4.1 Randomized Online Gradient Control (ROGC)

Let $\mathcal{J}_c \bigcup \mathcal{J}_d = \{1, \ldots, J\}$ denote the partition of the devices into convex and discrete ones, respectively. Note that for $j \in \mathcal{J}_d$, the feasible set $\mathcal{S}_n(j)$ is discrete, hence non-convex.

The ROGC algorithm is exactly the same as the OGC algorithm for the CC (cf. (2)) and every LC $j \in \mathcal{J}_c$ (cf. (3)). On the other hand, each LC $j \in \mathcal{J}_d$ performs the following:

(i) Advertise:

$$\mathcal{A}_{n+1}(j) := \Delta(\mathcal{S}_n(j)) \tag{12}$$

$$\hat{C}^{(j)}_{n+1}(y) := \sum_{s \in \mathcal{S}^{(j)}_n} C^{(j)}_n(s) y(s), \quad y \in \mathcal{A}_{n+1}(j), \tag{13}$$

where $\Delta(\mathcal{S})$ is the probability simplex imposed by a discrete set $\mathcal{S}$. Observe that $\mathcal{A}_{n+1}(j)$ is a convex set, and $\hat{C}^{(j)}_{n+1}(y)$ is the expected value of $C^{(j)}_n$ with respect to $y \in \Delta(\mathcal{S}^{(j)}_n)$, thus a linear function of y.

(ii) Compute:

$$y_n(j) = \mathcal{P}_{\Delta(\mathcal{S}_n(j))}(x_n(j)). \tag{14}$$

(iii) Implement *a random control* $S_n(j) \in \mathcal{S}_n(j)$ drawn from a probability distribution $y_n(j)$.

Let

$$R_n(\mathbf{z}_{1:n}) := \sum_{k=1}^{n} \left(\left[\sum_{j \in \mathcal{J}_c} w_j C_k^{(j)}(y_k(j)) \right] + \left[\sum_{j \in \mathcal{J}_d} w_j C_k^{(j)}(S_k(j)) \right] + G_k(y_k) - F_k(z_k) \right) \tag{15}$$

denote the dynamic regret of the ROGC algorithm at time step n with respect to $\mathbf{z}_{1:n}$ (cf. Definition 2). Note that $R_n(\mathbf{z}_{1:n})$ is a random variable due to appearance of the random variables $\mathcal{S}_n(j)$. The following result is a direct application of Theorem 1 to the ROGC algorithm.

Corollary 1. *Under the conditions of Theorem 1, the expected regret of the ROGC algorithm is bounded by*

$$\limsup_{n\to\infty} \mathbb{E}\left(\frac{R_n(\mathbf{z}_{1:n})}{n}\right) \leq K_1\alpha + \frac{K_2(1+\alpha\lambda)}{\alpha}\varepsilon + \frac{K_3}{\alpha}\limsup_{n\to\infty}\frac{V(\mathbf{z}_{1:n})}{n}$$

for any $\mathbf{z}_{1:\infty} \in \mathcal{X}_\infty$.

Remark 3. In our setting, we implicitly assume that the environment is *oblivious* in the sense that it produces the same sequence of objective functions and feasible sets regardless of the applied control actions. The high-probability regret bounds can thus be obtained similarly to that shown in [26, Lemma 1]. The extension to non-oblivious environments is a subject for future work.

Remark 4. Note that when the feasible set of a discrete device has large cardinality, the proposed ROGC algorithm might be impractical as it will require manipulations of large vectors. However, the ROGC algorithm can be extended to cover this case (or any case of *non-convex* bounded feasible set $\mathcal{S}_n(j)$) if instead of considering directly the probability simplex $\Delta(\mathcal{S}_n(j))$ as the advertised convex feasible set, one considers a convex hull $\mathsf{ch}(\mathcal{S}_n(j))$. The idea is to identify every $y \in \mathsf{ch}(\mathcal{S}_n(j))$ with a probability distribution $p_y \in \Delta(\mathcal{S}_n(j))$ parametrized by y such that $y = \mathbb{E}_{S\sim p_y}(S)$. The extension to this case is a subject of ongoing work.

4.2 Application to Real-Time Control of Electrical Grids

Consider the problem of controlling a collection of heterogeneous electrical resources that are interconnected via a portion of a power grid (e.g., a *distribution feeder* or a *microgrid*) under a typically time-varying objective and certain safety constraints. These resources can be photovoltaic (PV) systems, wind power plants, batteries, buildings, and electric vehicles. The resources are typically connected to the network via power inverters, thus allowing for the direct control of the (active and reactive) power setpoints at the point of connection. This problem has recently received renewed interest through the advent of renewable energy sources, such as solar power, and improved battery and inverter technologies.

The network-wide objective of the CC is to keep the power grid within the operational constraints (e.g., keeping node voltages and line currents within limits). Another possible goal of the CC is to ensure that the power flow at the point of connection to the higher-level grid follows a given time-varying signal – namely, making this subgrid *dispatchable*. We next outline a concrete real-time control problem in the spirit of [12, 5].

Consider an electrical distribution system comprising $J+1$ nodes collected in the set $\mathcal{J} \cup \{0\}$, $\mathcal{J} := \{1, \ldots, J\}$. Node 0 is defined to be the distribution substation at which the voltage is fixed. Let $V_j \in \mathbb{C}$ denote the voltage phasor at node $j = 1, \ldots J$, and let $v := [|V_1|, \ldots, |V_J|]^\mathsf{T} \in \mathbb{R}^J$

denote the vector collecting the voltage magnitudes. Without the loss of generality, we assume that there is a resource connected at every node $j \in \mathcal{J}$, thus identifying node j with LC j. The control variable for each resource j is given by $x(j) = (P(j), Q(j))^\mathsf{T} \in \mathbb{R}^2$, where $P(J)$ and $Q(j)$ are the active and reactive power setpoints, respectively. As a convention, positive power means production, whereas negative power signifies consumption. Let $x = (x(1)^\mathsf{T}, \dots, x(J)^\mathsf{T})^\mathsf{T} \in \mathbb{R}^{2J}$ collect the setpoints of the J resources. The relationship between v and x is given by the well-known nonlinear alternating-current (AC) power-flow equations $f(v, x) = 0$ (see, e.g., [3]).

To design the controllers, the nonlinear power-flow equations f are typically convexified or linearized around the current operation point. For the purpose of the example here, consider a possibly time-varying linear approximation to f in the form

$$v = A_n x + a_n, \tag{16}$$

where the system-dependent matrix $A_n \in \mathbb{R}^{2J \times J}$ and vector $a_n \in \mathbb{R}^J$ can be computed in a variety of ways (e.g., [2, 10, 13, 14, 7] and pertinent references therein). Similarly, the active power flow at the substation $P(0)$ (namely, the power that is exported to the higher-level grid) can be approximated as

$$P(0) = w_n^\mathsf{T} x + b_n, \tag{17}$$

for some $w_n \in \mathbb{R}^{2J}$ and $b_n \in \mathbb{R}$.

4.2.1 Design of the Central Controller

The CC obtains the advertisements $(\mathcal{A}_{n+1}(j), \hat{C}_{n+1}(j))$ from LCs $j \in \mathcal{J}$; see Section 4.2.2 below for details on how the LCs compute those. It also receives the estimation of the implemented power setpoint $\hat{y}_n$. It then constructs the objective function F_n according to (1). The network-wide objective G_n is designed using (17) to track a given sequence of power setpoints $\{P_{0,n}^{set}\}$ at the substation:

$$G_n(x) = 0.5 \left(w_n^\mathsf{T} x + b_n - P_{0,n}^{set} \right)^2. \tag{18}$$

The network-wide feasibility constraints are constructed using (16) as

$$\mathcal{U}_n := \{x \in \mathcal{A}_{n+1} : V_{\min} \leq (A_n x + a_n)_j \leq V_{\max},\, j \in \mathcal{J}\} \tag{19}$$

which ensures that the voltage magnitudes are within the prescribed limits $V_{\min}$ and $V_{\max}$. Finally, the control variables for the next time step are computed using (2).

4.2.2 Design of Local Controllers

For the purpose of this example, suppose that every resource is either (i) a PV system; (ii) a heating, ventilating, and air-conditioning (HVAC) system; or (iii) a battery. These three types of devices cover most modern distributed energy resources. Indeed, the PV system represents a volatile renewable power generator, the HVAC system represents a non-convex (discrete) controllable load, and the battery represents a bi-directional energy-storage resource. For every type of resource, we next present typical cost functions and feasibility constraints that are used to construct the advertisements to the CC.

For a PV system with inverter's rated power $S_{\text{inv}}(j)$ and an available active power $P_{\text{av},n}(j)$, the set $\mathcal{S}_n(j)$ is given by

$$\mathcal{S}_n(j) = \left\{ (P(j), Q(j)) \colon\, 0 \leq P(j) \leq P_{\text{av},n}(j),\, Q(j)^2 \leq S_{\text{inv}}(j)^2 - P(j)^2 \right\};$$

see, e.g., [25, 12]. Note that for PV inverters, the set $\mathcal{S}_n(j)$ is convex, compact, and time varying (it depends on the available power $P_{\text{av},n}(j)$, which in turn depends on the solar irradiance). The associated cost function typically encourages active-power generation and penalizes reactive power, e.g.,

$$C_n^{(j)}(P(j), Q(j)) = -c_1 P(j) + c_2 Q(j)^2 \tag{20}$$

for some positive constants c_1 and c_2. Note that a PV system is a convex resource in the terminology of Section 4.1, and thus the advertisement is defined as $\mathcal{A}_{n+1}(j) = \mathcal{S}_n(j)$ and $\hat{C}_{n+1}^{(j)} = C_n^{(j)}$.

Consider now a simplified case of an HVAC system that can be in either the ON or OFF state. When the system is in the ON state, it consumes $P_{\max}(j)$ active power and 0 reactive power. Moreover, the system can be *locked* in either the ON or OFF state because of cycling limitations and other constraints. Let $\ell_n(j) \in \{0, 1\}$ denote a binary state variable that equals 1 if the system is locked at time step n. The details of the related state machine are omitted as they are not essential for this example. With this at hand, the finite set of feasible setpoints $\mathcal{S}_n(j)$ is given by

$$\mathcal{S}_n(j) = \begin{cases} \{0, -P_{\max}(j)\}, & \text{if } \ell_n(j) = 0, \\ \{0\}, & \text{if } \ell_n(j) = 1,\ P_{n-1}(j) = 0 \\ \{-P_{\max}(j)\}, & \text{if } \ell_n(j) = 1,\ P_{n-1}(j) = -P_{\max}(j). \end{cases}$$

Finally, the cost of being in the ON state, $C_n^{(j)}(-P_{\max}(j))$, and in the OFF state, $C_n^{(j)}(0)$, is system-dependent and reflects, for example, the current temperature and its distance from the desired setting. Observe that an HVAC system is a discrete resource in the terminology of Section 4.1, and therefore the advertisement is defined by (12) and (13). In particular, let $y \in [0, 1]$ denote the probability to turn the HVAC system on. Then

$$\mathcal{A}_{n+1}(j) = \begin{cases} [0, 1], & \text{if } \ell_n(j) = 0, \\ \{0\}, & \text{if } \ell_n(j) = 1,\ P_{n-1}(j) = 0 \\ \{1\}, & \text{if } \ell_n(j) = 1,\ P_{n-1}(j) = -P_{\max}(j) \end{cases}$$

and for $y \in \mathcal{A}_{n+1}(j)$, $\hat{C}_{n+1}^{(j)}(y) = (1 - y)C_n^{(j)}(0) + yC_n^{(j)}(-P_{\max}(j))$.

Finally, consider a battery with state-of-charge at time step n given by $\text{SoC}_n \in [0, 1]$. Let $P_{\min,n}(j)$ and $P_{\max,n}(j)$ denote, respectively, the lower and upper bounds on active power production. These are time-varying quantities that depend on operating conditions of the battery, such as SoC_n and the DC voltage; see, e.g., [25]. With inverter's rated power $S_{\text{inv}}(j)$, the set $\mathcal{S}_n(j)$ is given by

$$\mathcal{S}_n(j) = \left\{(P(j), Q(j)):\ P_{\min,n}(j) \le P(j) \le P_{\max,n}(j),\ Q(j)^2 \le S_{\text{inv}}(j)^2 - P(j)^2\right\}$$

similarly to the PV system. The associated cost function can be designed based on the SoC_n and the desired value for the state-of-charge. For example, if SoC_n is greater than the desired value, a function that encourages power production can be defined, similarly to (20). Conversely, if SoC_n is smaller than the desired value, a function that encourages power consumption can be defined, e.g.:

$$C_n^{(j)}(P(j), Q(j)) = c_1 P(j) + c_2 Q(j)^2$$

for positive c_1 and c_2. Note that, similarly to the PV system, a battery is a convex resource, thus $\mathcal{A}_{n+1}(j) = \mathcal{S}_n(j)$ and $\hat{C}_{n+1}^{(j)} = C_n^{(j)}$.

To conclude this section, we note that in [5] a special case of the OGC algorithm was used as an heuristic to control a realistic power system (a microgrid), and it was shown that this algorithm performed well numerically.

5 Conclusion and Future Work

We presented and analyzed a first-order control algorithm that bridges the gap between online convex optimization and real-time control. We showed that this algorithm possesses small dynamic regret under certain conditions on the measurement error and time variability of the optimal trajectory. The algorithm can be applied to control heterogeneous resources in real time, and in particular to control a mix of convex and discrete resources. We also illustrated the application of the algorithm in the context of real-time control of the electrical grid.

The proposed OGC algorithm is only a first (and the most straightforward) example of control algorithms that can be applied in the proposed framework. An extension to other methods to optimize the setpoints is of interest. In particular, it is an interesting question whether a distributed algorithm (e.g., based on primal-dual decomposition method as in [12] or on the alternating direction method of multipliers (ADMM)) can be analyzed similarly to show the small dynamic regret. Further, extending the framework to general non-convex feasible sets seems possible (cf. Remark 4) and is a subject of ongoing work. Finally, an interesting research direction is in devising online control algorithms based on the concept of *approachability* [6], which is a more general concept than no-regret algorithms.

References

[1] D. Adamskiy, W. M. Koolen, A. Chernov, and V. Vovk. A closer look at adaptive regret. In *Proceedings of the 23rd International Conference on Algorithmic Learning Theory*, ALT'12, 2012.

[2] M. E. Baran and F. F. Wu. Network reconfiguration in distribution systems for loss reduction and load balancing. *IEEE Trans. on Power Delivery*, 4(2):1401–1407, Apr. 1989.

[3] A. R. Bergen and V. Vittal. *Power System Analysis*. 2nd ed., Upper Saddle River, N.J. : Prentice Hall, 2000.

[4] A. Bernstein, N. J. Bouman, and J.-Y. Le Boudec. Real-time minimization of average error in the presence of uncertainty and convexification of feasible sets, 2016. arXiv:1612.07287.

[5] A. Bernstein, L. E. Reyes-Chamorro, J.-Y. Le Boudec, and M. Paolone. A composable method for real-time control of active distribution networks with explicit power set points. Part I: Framework. *Electric Power Systems Research*, 125:254–264, August 2015.

[6] D. Blackwell. Controlled random walks. In *Proceedings of the International Congress of Mathematicians*, volume III, pages 335–338, 1954.

[7] S. Bolognani and F. Dörfler. Fast power system analysis via implicit linearization of the power flow manifold. In *2015 53rd Annual Allerton Conf. on Communication, Control, and Computing*, pages 402–409, 2015.

[8] M. Brown, J. Funke, S. Erlien, and J. C. Gerdes. Safe driving envelopes for path tracking in autonomous vehicles. *Control Engineering Practice*, 2016. To appear.

[9] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA, 2006.

[10] K. Christakou, J.-Y. Le Boudec, M. Paolone, and D.-C. Tomozei. Efficient Computation of Sensitivity Coefficients of Node Voltages and Line Currents in Unbalanced Radial Electrical Distribution Networks. *IEEE Transactions on Smart Grid*, 4(2):741–750, 2013.

[11] E. Dall'Anese, S. V. Dhople, and G. B. Giannakis. Optimal dispatch of photovoltaic inverters in residential distribution systems. *IEEE Trans. Sust. Energy*, 5(2):487–497, Apr. 2014.

[12] E. Dall'Anese and A. Simonetto. Optimal power flow pursuit. *IEEE Transactions on Smart Grid*, 2016. To appear.

[13] S. Dhople, S. Guggilam, and Y. Chen. Linear approximations to AC power flow in rectangular coordinates. *Allerton Conference on Communication, Control, and Computing*, 2015.

[14] S. Guggilam, E. Dall'Anese, Y. Chen, S. Dhople, and G. B. Giannakis. Scalable optimization methods for distribution networks with high PV integration. *IEEE Transactions on Smart Grid*, 2016.

[15] E. C. Hall and R. M. Willett. Online convex optimization in dynamic environments. *IEEE Journal of Selected Topics in Signal Processing*, 9(4):647–662, June 2015.

[16] E. Hazan. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.

[17] E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2):169–192, 2007.

[18] E. Hazan and C. Seshadhri. Efficient learning algorithms for changing environments. In *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML '09, pages 393–400, 2009.

[19] M. Herbster and M. K. Warmuth. Tracking the best linear predictor. *J. Mach. Learn. Res.*, 1:281–309, September 2001.

[20] P. Jahangiri and D. C. Aliprantis. Distributed Volt/VAr control by PV inverters. *IEEE Transactions on Power Systems*, 28(3):3429–3439, Aug 2013.

[21] G. Neu and M. Valko. Online combinatorial optimization with stochastic decision sets and adversarial losses. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems (NIPS) 27*, pages 2780–2788. 2014.

[22] S. Paudyal, C. A. Canizares, and K. Bhattacharya. Optimal operation of distribution feeders in smart grids. *IEEE Trans. on Ind. Electron.*, 58(10):4495–4503, Oct. 2011.

[23] R. Potluri and A. K. Singh. Path-tracking control of an autonomous 4WS4WD electric vehicle using its natural feedback loops. In *2013 IEEE International Conference on Control Applications (CCA)*, pages 394–400, Aug 2013.

[24] G. V. Raffo, G. K. Gomes, J. E. Normey-Rico, C. R. Kelber, and L. B. Becker. A predictive controller for autonomous vehicle path tracking. *IEEE Transactions on Intelligent Transportation Systems*, 10(1):92–102, March 2009.

[25] L. E. Reyes-Chamorro, A. Bernstein, J.-Y. Le Boudec, and M. Paolone. A composable method for real-time control of active distribution networks with explicit power set points. Part II: Implementation and validation. *Electric Power Systems Research*, 125:265–280, August 2015.

[26] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the Twentieth International Conference on Machine Learning, (ICML 2003), August 21-24, 2003, Washington, DC, USA*, pages 928–936, 2003.